

JPL Publication 10-15



Vehicle-Terrain Interaction Modeling and Validation for Planetary Rovers

Dartslab (<https://dartslab.jpl.nasa.gov/>)

Jonathan Cameron
Abhi Jain
Terry Huntsberger
Garett Sohl
Rudranarayan Mukherjee

**National Aeronautics and
Space Administration**

**Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California**

August 2009

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

©2009 California Institute of Technology. Government sponsorship acknowledged.

Vehicle-Terrain Interaction Modeling and Validation for Planetary Rovers*

Jonathan Cameron

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91107
jonathan.m.cameron@jpl.nasa.gov

Abhi Jain

Jet Propulsion Laboratory
California Institute of Technology
abhi.jain@jpl.nasa.gov

Terry Huntsberger

Jet Propulsion Laboratory
California Institute of Technology
terry.huntsberger@jpl.nasa.gov

Garett Sohl

Jet Propulsion Laboratory
California Institute of Technology
garett.sohl@jpl.nasa.gov

Rudranarayan Mukherjee

Jet Propulsion Laboratory
California Institute of Technology
rudranarayan.mukherjee@jpl.nasa.gov

Abstract

Physics-based modeling and simulation of planetary rovers is an integral part of planning, testing and design of robotic planetary missions such as the upcoming Mars Science Laboratory (MSL) and the extremely successful Mars Exploration Rovers (MER) Spirit and Opportunity. As planetary rover missions grow more challenging with heavier rovers moving at higher speeds over rougher terrain than in the past, modeling the vehicle terrain interactions has become a critical part of the success of model-based testing and operations within the NASA / JPL physics-based ROAMS rover simulator. This paper focusses on the wheel-terrain contact model used in the ROAMS simulator and its validation for the FIDO class rovers. It also presents a brief overview of modeling planetary rovers and terrains within the ROAMS simulator.

Since the ROAMS simulator is used in various modes such as stand-alone simulation, closed-loop simulation with on-board software simulation, or for operator-in-the-loop simulations, faster than real time computational performance is an essential requirement. Consequently, the vehicle terrain interaction model is based on a physics approach designed to retain adequate fidelity without compromising performance while providing the ability to vary parameters to model different soil and rover properties.

*Copyright 2009 California Institute of Technology. Government sponsorship acknowledged.

1 Introduction

The exploration of Mars by the twin Mars Exploration Rovers (MER) Spirit and Opportunity has led to exciting discoveries and has been an inspiration to many. The MER rover Spirit has traversed over 7.7 kilometers and Opportunity has traversed more than 17 kilometers since landing on Mars in early 2004, far exceeding the mission requirements. Some of those traverses were conducted in regions such as hillsides or crater walls where the rover wheels experienced significant slipping on the terrain. The safety and success of such challenging driving has been largely because of empirical methods and operational rules specific to the MER rover that are used to minimize and control slip. The methods and rules have been learned mostly from experimental vehicle tests here on the Earth.

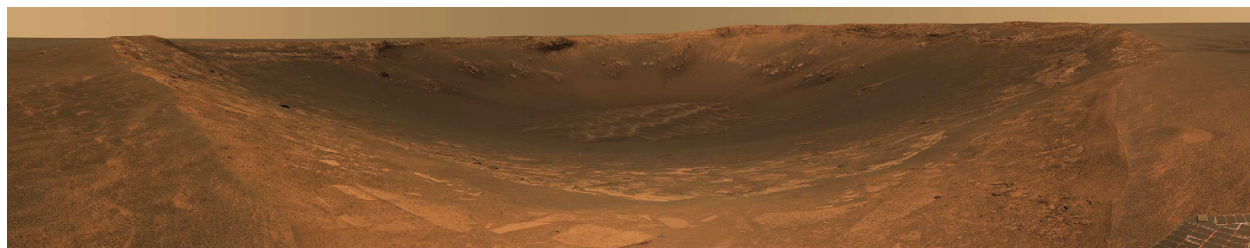


Figure 1: Endurance Crater on Mars. The MER rover Opportunity drove through many challenging areas in Endurance Crater. The safety of the rovers and the success of the MER mission depends largely on good vehicle modeling and operational rules based on models and experimental results.

Testing driving sequences as well as creating and validating conservative operational rules based on experimental vehicle tests is expensive and very limited in the range of scenarios that can be analyzed. In the next mission to Mars, the Mars Science Laboratory (MSL) rover, will driver much further over rougher terrain. As mission requirements grow for driving longer traverses, traversing rougher terrain, and requiring more autonomy, there is a significant need to expand the scope of the vehicle drive planning design, test, and verification process. A tool for examining slip behavior of a general class of rovers over more widely variable terrain is needed for upcoming missions.

Currently there is a lack of an integrated capability that can fill the need for an integrated model that can truly serve as a surrogate model. NASA and JPL set out to develop the **Rover Analysis, Modeling and Simulation (ROAMS)** planetary rover modeling capability to fill this gap. The key requirement driving the ROAMS simulator development are:

- sufficient fidelity in the model physics so that its behavior matches the real system,
- adequate speed so it can be used with flight software,
- detailed low-level interfaces and models so it can close the loop with flight software,
- full model suite so all loops with the onboard software can be closed,
- validation experiments to verify the simulation physics,

- realistic planetary environment models (geometry, soil mechanics) so it can expand scenarios to study, and
- multi-mission capability so it can be extended to different configurations.

The development, design and validation of ROAMS is driven by the unique challenges of semi-autonomous design and operation of planetary rovers in remote environments. Unlike commercial/military terrestrial vehicles, the amount of data available about the terrain environment is mostly unknown and meagre when available. At best, the information about the terrain and rover interaction properties are refined and improved from the downlink telemetry acquired from the traverse history. Indeed, the current state of practice in MER operations is limited to using augmented kinematic models to predict rover slippage over terrains during rover operations (Yen, 2008). This is in contrast with the large number of well instrumented, controlled experiments that can be done in terrestrial applications to improve the fidelity of simulation models. Flight rovers have to be designed, tested and validated to operate in areas that have never been visited before.

These circumstances pose different priorities and different requirements on the vehicle modeling strategy for space robotics than for typical terrestrial applications. The new dynamics models represent a significant advancement over the prior kinematic approaches. The goal is not to build a single high fidelity model of the rover since there is not enough data to validate such a model. Further, the large tuning effort put into developing a single high-fidelity model is unlikely to remain usable across the range of environments that are typically encountered during traverses. Instead, the expectation on ROAMS is to provide a parametrized model whose parameters can be tuned based on available data, so that it can be used as a predictive model during operations. In an operational context, the simulator parameters are constantly tuned and updated to match the current environmental context for the rover. The validation requirement on ROAMS is to show that the model generalizes to handle the full range of motions (cross-slope, up-hill, down-hill, turns, *etc.*) typical in traverses in a natural environment. This is a strategy adopted in the ROAMS validation approach used in this paper, where we have observed that a small subset of experimental data can be used to tune the model to predict performance over a wide range of different motions across different terrain geometries. This approach is used in both terrestrial ground tests of the rover as well as during rover operations.

In addition model fidelity, equally challenging requirements on ROAMS arise from its use in hardware-in-the-loop, closed-loop testing of the flight software. Thus ROAMS is used in ground rover testbeds to provide a high-fidelity model that closes all loops (motors, IMUs, cameras) with the flight software in real-time environments. These same testbeds support rover operations once the rovers are deployed. The fidelity requirements on ROAMS are driven by the need to provide realistic behaviour that can be used to develop and test flight algorithms under a variety of environments and conditions that would be otherwise difficult through physical tests. Indeed, the ROAMS-based tests complement “hardware” simulation tests done in the field with engineering prototypes of the flight rovers. The faster than real-time performance requirements on ROAMS have led to the use of high performance algorithms and modeling techniques, as well as engineering judgements to judiciously design models that provided adequate fidelity without compromising performance speed. To get a

sense of the model complexity, the model for just the MSL 5 degrees of freedom arm alone has over 17 degrees of freedom to adequately model the motors, backlash and joint flexibility effects essential to adequately testing the flight algorithms and software. Thus ROAMS represents a comprehensive, end-to-end simulator where the sinkage/slippage models play a critical role among a larger class of requirements.

This paper provides an overview of the ROAMS design and architecture in Section 2, followed by a brief description of the vehicle and terrain modeling approaches. Section 3 describes the vehicle/terrain interaction models in detail largely based on (Sohl and Jain, 2005). Section 4 details the model validation against MER data, and Section 5 describes validation experiments and analysis by the Mars Technology Program based on (Huntsberger et al., 2008).

2 Overview of the ROAMS planetary rover simulator

In the following sections we give some background on the design and goals of the ROAMS infrastructure.

A primary requirement on ROAMS is that it serve as a high-fidelity surrogate rover to support closed-loop testing beyond what is possible with just hardware rover testbeds. These high fidelity needs require ROAMS to implement (a) detailed physics-based models of the rover mechanical platform including its kinematics and dynamics, (b) its suite of actuators and sensors such as wheel and steering motors and encoders, inertial measurement units (IMUs), sun sensors, cameras, and (c) models of the environment and the rover's interactions with the environment. Hand in hand with the model development process is an ongoing ROAMS simulator validation effort consisting of a series of experiments involving deterministic as well as statistical comparisons with physical rover data.

Development of the rover flight system typically involves test platforms ranging from experimental technology development rovers all the way to flight breadboards and spares. The configuration of these platforms typically evolves over time with updates to the sensor/actuator suite, avionics and other hardware components. ROAMS is designed to provide models that shadow these multiple rover platform configurations at any given time and track their evolution over time. This required that ROAMS avoid monolithic, rover platform specific simulation implementations. Instead a conscious design strategy has been to allow users to configure ROAMS for different rover models easily at run-time via model data files. While allowing users to easily tailor simulations to the specific platforms, this configurability has been useful during the simulation validation effort to match ROAMS to rover model configurations used in the experiments.

As a test platform, ROAMS is used in closed-loop with the on-board rover software and hardware. This requires ROAMS to be embeddable within closed-loop testbed environments containing a mix of on-board software, real hardware and simulated hardware. ROAMS provides hardware-like command and sensing interfaces similar to actual hardware to allow such loop closure. Particular attention has been paid to simulation algorithm performance in order to meet the closed-loop timing requirements. Also, ROAMS is portable across Unix

and real-time VxWorks platforms.

While simulations are expected to do the “right” thing, i.e. provide good fidelity, they also need to provide a significant level of instrumentation and other features for them to be usable. Since the inclusion of these features adds to code size and the number of external dependencies, ROAMS has adopted a layered design, where many of the features are implemented as optional plug-in extensions so they can be included as needed at run-time. This approach has also helped increase the amount of reusable modules within ROAMS.

ROAMS is built upon the existing DARTS and Dshell simulation framework (Biesiadecki et al., 1999) developed for spacecraft simulations. The ROAMS development has focused primarily on model development needed for the surface rover domain. While the ROAMS core is implemented in C/C++. It includes a Python scripting interface (auto-generated by the SWIG wrapper generation tool) to the core C/C++ classes to facilitate simulation configuration and regression testing. This scripting capability is also used to develop graphical user interfaces for users to change simulation modes, set rover goals, change simulation speed, take time steps, exercise rover degrees of freedom, select terrain models, *etc.* The Dspace 3D visualization tool (Jain et al., 2003) provides run-time visualization of the rover simulation state.

More details on ROAMS are available in (Jain and Yen, 1999) (Jain et al., 2003) (Jain et al., 2004).

2.1 Vehicle models

The mechanical aspects of the vehicle are modeled using $O(n)$ highly efficient recursive multi-body dynamics algorithms framework, called the Spatial Operator Algebra, developed over the last two decades at JPL (Rodriguez et al., 1992). The underlying efficiency in the algorithms design enables high computational efficiency in modeling the kinematics and dynamics of the vehicles. The kinematics and dynamics of various components such as steered wheels, suspensions of various kinds, chassis dynamics, mast or arm motion, are accurately and efficiently modeled with faster than real time performance. The kinematics and dynamics of prescribed (controlled) and free (uncontrolled) joint motions are seamlessly captured in a unifying mathematical framework. Both joint accelerations and joint constraint forces can be calculated using these algorithms. The external or active forces and torques acting on the vehicles include gravity, terrain interaction loads, and motor torques.

The vehicles feature physics-based actuator and sensor models. The wheel and steer motors are modeled as DC motors with damping and back EMF. The models are parametric in nature and different motors can be modeled by varying the motor characteristics such as armature resistance, current and torque constants. The sensor models are typically used to log the states of the system through IMUs and encoders. Full 3D visualization of the vehicles is enabled through VRML-based part graphics of various bodies in the system. The part graphics are articulated based on the kinematic model of the system, enabling realistic representation of the vehicle motion. Enhancements on the vehicle models also include various additional physics-based models such as thermal models, power and battery models,

models for charging from solar panels, and life support systems models.

Figure 2 illustrates a block diagram of the parametric vehicle model scheme. The core of the vehicle is based on the dynamics methods, sensor and actuator models. ROAMS contains implementations of controller models for individual motor control. There are also hooks in place to enable integration with an external motor controller or hardware. Similarly ROAMS contains navigation and collision detection algorithms for autonomous vehicle mobility and manipulation. There are nonlinear camera models that enable visual odometry and range mapping. ROAMS contains models of various artificial and natural terrain as well as algorithms for modeling the interaction of the vehicle with the terrain. This is discussed in detail in the following sections.

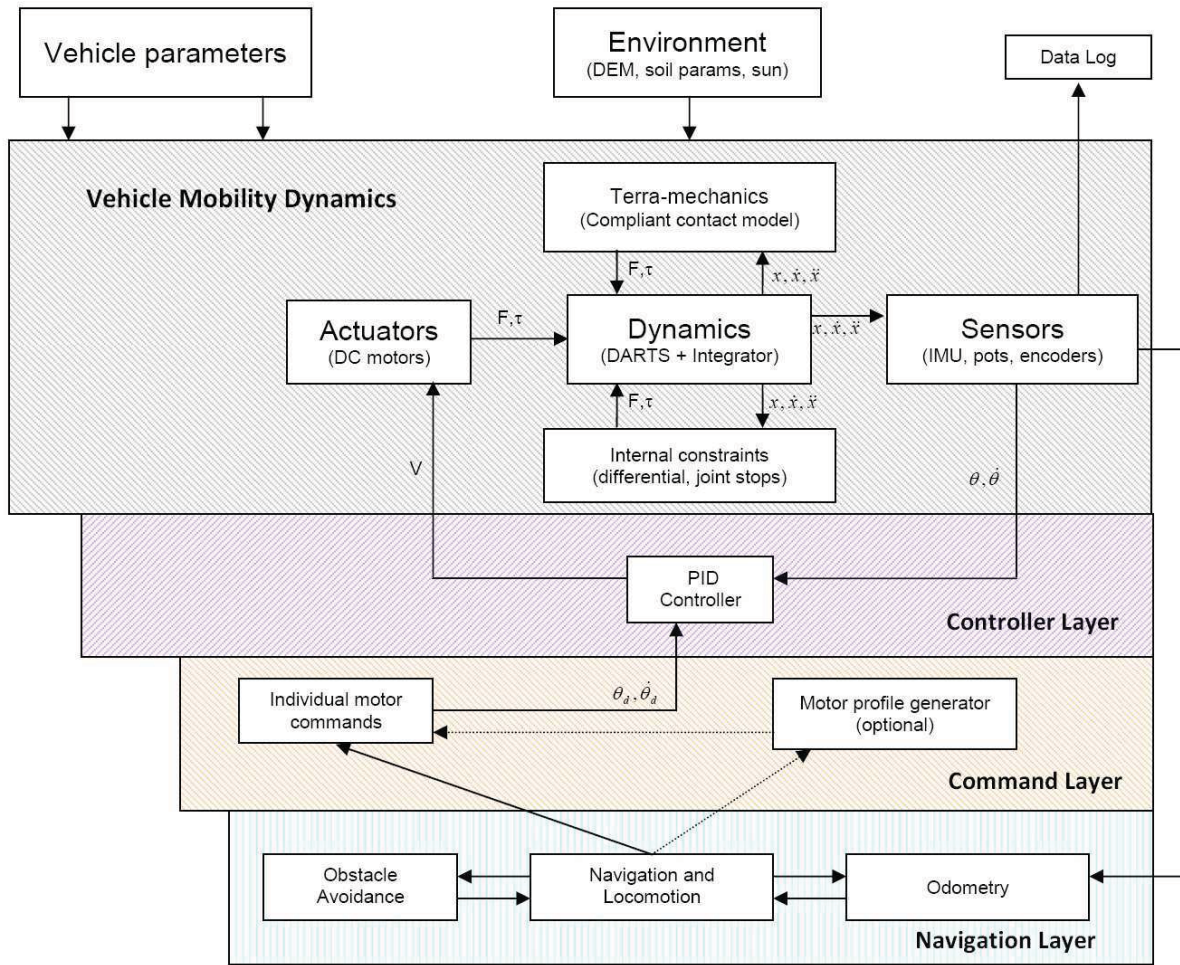


Figure 2: Overview of ROAMS software. The mobility includes vehicle dynamics, sensor, actuator, and environmental models.

2.2 Environmental and terrain models

As a surface vehicle, the wheels of the rover interact with the terrain environment. ROAMS uses the SimScape terrain modeling toolkit (Jain et al., 2006) for its terrain models. A Digital Elevation Map (DEM) represents the terrain (including rocks) in ROAMS. ROAMS parameterizes soil mechanics properties using density, internal friction angle and soil cohesion. The user may explicitly set these parameters or select parameter sets that match certain soil types (clay, loose sand, mixed, *etc.*). ROAMS needs to incorporate raw terrain data from a variety of sources including planetary data archives, field tests, synthetic models (derived from existing terrain models) as well as algorithmically generated ones. Such primary terrain data needs to be processed and assembled into a terrain model prior to use by applications. This often requires piecing together data from different sources, processing relative spatial transforms among them, assigning attributes to the geometry, handling different data formats and representations and using height field and/or 3D models as appropriate

SimScape supports multiple representations of the terrain geometry. These include 2.5D digital elevation map grid representations, point cloud representations, 3D mesh representations, and 2.5D triangulated irregular network (TIN) mesh representations and support for geo-referenced planetary data models. It also includes algorithms and methods for transformations between different terrain model representations. SimScape can be used as an embedded library with efficient algorithms and methods for interacting with terrain models. Terrain model data from sources such as planetary archives, field test sites, and synthetic and analytic terrain models can be used within ROAMS. There is support for exporting and importing terrain models to and from a variety of standard terrain data formats including PDS, GeoTiff, USGS ISIS, *etc.* Furthermore, there is support for overlaying of surface properties such as material composition, texture, albedo, terra-mechanics parameters, reflectivity, and user defined properties onto the underlying terrain geometry for run-time use by simulation models. Terrain models can be assembled and stored for later retrieval and use by the ROAMS simulations.

3 ROAMS wheel-terrain contact model

The goal of the wheel-terrain contact model is the determination of the contact forces and torques exerted by the terrain on the rover wheels. These contact forces, along with wheel motor torques, provide the motive force for the rover and allow it to traverse over the terrain. The forces and torques must also excite on-board sensor models, such as gyros and accelerometers in a realistic fashion.

Before considering the details of the model in ROAMS, a brief overview of the prevalent modeling techniques is presented.

3.1 Current approaches to vehicle-terrain interaction modeling

The most widely known methods for vehicle terrain interactions modeling are the Bekkers Derived Terramechanics Model (BDTM) (Bekker, 1956)-(Bekker, 1960)-(Bekker, 1969) and

the Cone Penetrometer Technique (Wong, 2001). In the 1960's Bekker developed the Bekkers Derived Terramechanics Model (BDTM) based on a combination of scientific laws and empirical data. This model captures pressure-sinkage relationships as well as slip-shear phenomena through the use of various model parameters. On the other hand, the Cone Penetrometer Technique uses a single parameter called Cone Index to characterize the soil strength based on experimental measurements of penetration and normal force in cone penetrometer tests. Based on either of these methods, references have been developed for measuring vehicle mobility through vehicle terrain interactions, for example (Haley et al., 1979a)-(Haley et al., 1979b)-(Ahlvin and Haley, 1992). There is an exhaustive volume of work related to this topic and representative methods are discussed below.

Broadly classified, there are three main methods that have traditionally been used to model interactions between vehicles with terrains. These interactions are not just limited to wheel soil interactions, but include terrain modifications such as excavation. The first of these methods is based on continuum mechanics based approaches that use finite element methods or discrete element methods to model the interactions, examples of which are (Singh, 1997)-(Solisa and Longoria, 2008). These methods model the deformation of the terrain in detail and can be used to develop high fidelity models. However, these methods are relatively slow and cannot be used in real-time simulations or closed-loop validation.

The second set of methods based on the first principles approach motivated by the seminal work of Terzaghi (Terzaghi, 1943) and others. These methods (D. J. van Wyk, 1996)-(Q. Li, 2007)-(Ellery, 2005)-(Andrade et al., 1998)-(Hutangkabodee et al., 2008) use friction, part geometry and other physical properties to predict approximate interaction forces.

The third approach is based on engineering mechanics approaches where simplified models are used to approximate the interaction forces between the terrain and mechanical parts without developing high fidelity soil models (Wong, 2001)-(Iagnemma et al., 1999)-(Wong and Huang, 2006)-(Ward and Iagnemma, 2007)-(Shibly et al., 2005). These methods typically employ an analytical deformation model similar to a spring damper excitation to model the interactions and forces.

A new trend in modeling terrain vehicle interaction is through the use of stochastic methods and approximation of uncertainties using polynomial chaos (Sandu et al., 2006)-(Kewlani and Iagnemma, 2008). Depending on the research, the methods usually apply uncertainty in the terrain heights, introduce uncertainty in terrain properties such as friction and cohesion or introduce uncertainty directly in experimentally generated parameters for analytical models.

The automotive industry uses different tire models to model the interactions between the road and the vehicle. Most notable among them are the Pacejka (Pacejka and Bakker, 1992) formulae. Custom tire models have also been proposed by research (Chan et al., 2005) or commercial proprietary simulation packages such as ADAMS (ADAMS, 2000) in the realm of vehicle dynamics both for on and off-road performance. These models are most suitable for vehicles with rubber or deformable tires. and are significantly different from rover dynamics where the wheels are modeled as rigid and speeds are significantly slower.

Several software packages are also available for planetary rover simulation and analysis.

Notable among the commercially available packages is ADAMS (ADAMS, 2000) while examples of research codes include (Bauer et al., 2005)-(Thueer et al., 2007)-(Poulakis et al., 2008)-(Jain et al., 2003).

3.2 Statically indeterminate techniques

In a general formulation, there are three unknown force and three unknown torque components that define the net effect of each contact on the rover wheel. For a six wheeled rover, there are a total of 36 unknowns. However, a static force analysis of the rover provides only 6 equations (three linear and three angular) in these 36 unknowns. For the specific case of a rocker-bogey mechanism, three additional constraint equations can be generated (one rocker differential and two bogey constraints). However, even with those added moment constraints, the problem is still statically indeterminate: 9 equations in 36 unknowns.

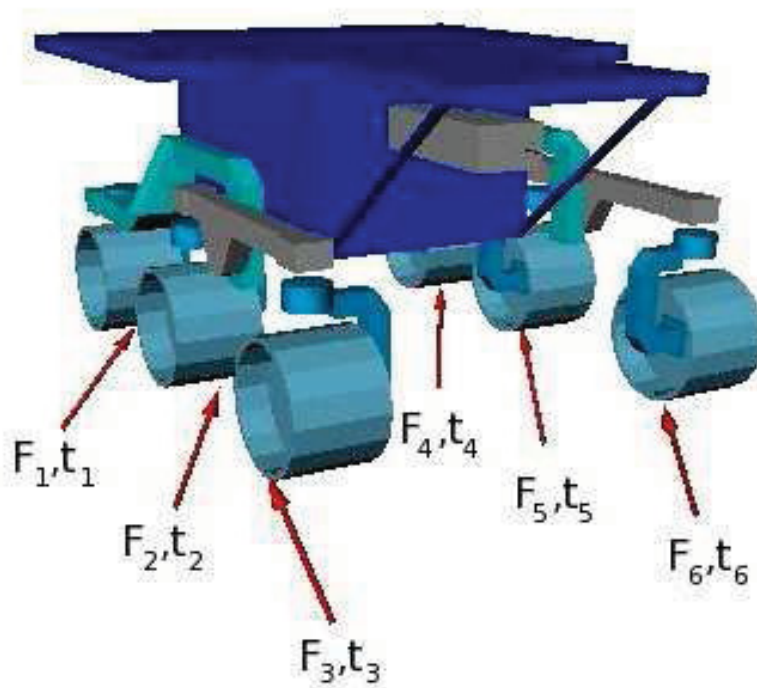


Figure 3: Rover contact forces and torques

The contact forces and torques on the six wheels of a rover in contact with the terrain are statically indeterminate. Sometimes it is possible to simplify the problem by eliminating unknowns until the problem can be solved. Many contact model formulations assume that the soil can only exert moments on the wheel about the terrain normal direction. This eliminates two unknowns for each wheel (reducing the total number of unknowns to 24). It is also possible to assume an effective point-plane contact between a point on the wheel and the “plane” of the soil. In this case, all the moments are zero (reducing the total number of unknowns to 18) (Hacot et al., 1998). ROAMS uses the point-plane contact assumption

in determining the contact forces. ROAMS also estimates a surface contact area on each wheel based on wheel sinkage. This contact area is used in determining maximum allowable traction 2.

Another common assumption used to reduce the complexity of the problem is to assume that the vehicle roll angle is small (Iagnemma et al., 1999). Under this assumption, the transverse force component is zero for each wheel contact. This assumption is often used to decompose the six wheel rover into two planar problems (Farritor et al., 1998). Each planar problem contains three wheels on the left or right side of the rover. The planar problem is then analyzed in detail. This assumption, however, is not suitable to a general dynamics simulation since transverse forces can not be assumed to be zero for motions over undulating terrain.

Some approaches attempt to compute contact forces satisfying a set of validity constraints (Hait and Siméon, 1996). For example, contact forces that fall within the friction cones at each contact point and satisfy the static force/moment equations can be found. While these techniques can determine if a suitable set of contact forces exist, they do not guarantee uniqueness or continuity of the solution. In addition, they may fail to find solutions in cases that are outside the assumed solution space. This makes them very unsuitable for a general dynamics simulation where high, or even total, slippage conditions will be encountered.

Another approach is to solve for contact forces that optimize a given criteria (instead of merely satisfying a given criteria) (Hacot, 1998). Assuming a well posed problem and good optimization routines, a unique solution can be found. Continuity of solution as the rover moves over the terrain is not guaranteed, but can reasonably be expected in many cases. Criteria such as minimum energy are often used. This procedure is generally very time consuming and while the minimization of a single criteria produces a unique answer, there is no guarantee that answer is a correct solution for the contact forces. The time consuming nature of optimization techniques makes this method unsuitable for a real-time dynamics simulation where rapidly changing contact forces for each wheel must be computed hundreds of times per second.

3.3 Compliance-based techniques

Adding compliance to a statically indeterminate system can allow a solution to be found (Kraus et al., 1997) at the cost of adding system states and increasing numerical stiffness of the problem. The forces on an object as simple as a four legged table on a flat surface are statically indeterminate. However, the introduction of compliance in the four legs allows the contact forces to be directly computed. Adding compliance provides a unique solution that is generally correct under the limitations of the compliance model. However, allowable stiffness is often restricted by the choice of numerical integration algorithm (Alexander, 1994). If the actual stiffness can not be used (due to numerical stability problems) compliance techniques may only approximate the correct answer. However, it is observed that compliance techniques will degrade gracefully and still provide good approximations with less than ideal stiffness parameters.

Compliance techniques provide a solution for contact forces based on the deflection of a spring-damper system. This deflection is directly related to the state of the system. Since the system state is always available to a dynamics simulation, these computations can be done very quickly. This is highly desirable in a real-time simulation environment, where optimization or other search techniques are too slow.

The contact model in ROAMS assumes point-plane contact. The terrain under the wheel is assumed to be locally planar and the contact forces are applied to a single point on the wheel. Under this assumption, there are three unknown force components for each wheel contact. Two separate and independent compliance systems are used to compute the three force components. One compliance system is used to compute the force component in the normal direction. The normal direction is defined as perpendicular to the “plane” of the terrain. Once the normal force is computed, it can be used to estimate wheel sinkage (Bekker, 1969) and resulting contact area. Even though the force is applied at a single point on the wheel, the sinkage and resulting contact area is important in estimating maximum traction. The second compliance system is a two degree of freedom system. It is used to compute the two components of the force in the plane of the terrain. The next two sections will describe how the normal and tangent plane systems are implemented in ROAMS .

3.4 Normal force

The magnitude of the normal force is the foundation of almost every contact mechanics formulation. From a simple Coulomb friction model to a complex terra-mechanics model (Bekker, 1969), they all use the magnitude of the normal force to determine the available traction force in tangent directions. However, these formulations do not describe *how* to compute the normal force. It is assumed to be a given for the problem. In the case of a six wheeled rover in contact with undulating terrain (where the terrain normal at each contact point is different), these forces are statically indeterminate.

In order to compute the force in the normal direction, ROAMS uses a single degree of freedom, Hunt-Crossley (Hunt and Crossley, 1975) compliance system at each wheel:

$$F_N = k_N \delta_N^n + \frac{3}{2} \alpha_N \dot{\delta}_N \delta_N^n \quad (1)$$

where F_N is the force in the normal direction, k_N is a spring constant, α_N is a damping constant, n is the non-linear deflection exponent, and δ_N is the deflection. Figure 4 shows a side view of the normal direction compliance system.

The normal direction is determined by an examination of the terrain DEM for a small area under the wheel. There is an assumption that the terrain normal is relatively smooth and that the sinkage of the wheel into the terrain does not significantly effect the terrain normal at the contact point. The deflection of the compliance system is based on the penetration of the wheel into the terrain in the terrain normal direction. The location of the contact point on the wheel is defined as the point on the wheel that penetrates the farthest into the terrain. As the wheel penetrates the terrain, the compliance system applies an opposing force at the contact point. The opposing force will rapidly reach equilibrium with the weight of the rover supported by that wheel.

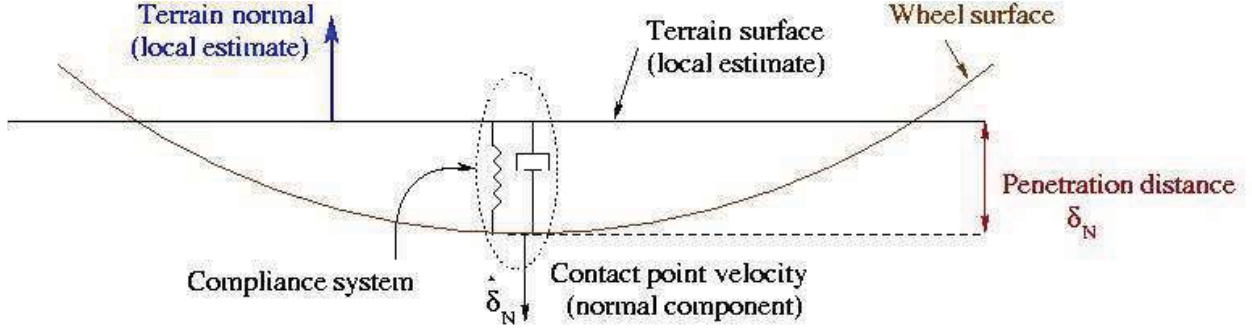


Figure 4: Normal direction compliance system

Equation 1 requires an estimate of the penetration of the wheel into the terrain. To simplify this computation, wheels are assumed to be cylindrical and the terrain under each wheel is assumed to be locally planar. These assumptions allow rapid computation of the penetration distance and contact point on the wheel based on the local terrain height field, wheel position and wheel attitude. Since the penetration of the wheel into the terrain is generally small and the weight supported by a given wheel may be substantial, ROAMS uses a very stiff spring in the normal direction (k_N is large) to prevent excessive penetration. This can lead to stability problems during numerical integration and often requires tuning the normal direction spring and damping constants based on the total rover mass. It is also important to note that the Hunt-Crossley compliance model in Equation 1 does not attempt to model the sinkage of a wheel into soil, but is rather an algorithmic convenience used to compute the statically indeterminate normal force. After a normal force has been found, it can be used to estimate the actual sinkage of a wheel into the soil based on soil properties such as density and cohesion. This results in two separate concepts of wheel-soil penetration: one from the compliance model and one from terra-mechanics.

3.5 Tangent plane forces

While the force in the normal direction provides “support” for the vehicle, it is the tangent plane forces that allow a wheeled vehicle to move. The tangent plane is defined as the plane perpendicular to the normal vector direction. Most contact models follow the basic premise that there is a limit to the magnitude of the tangent forces. Once that limit is reached, the tangent force can no longer prevent relative motion between the contact point and the terrain and the wheel starts to slip. The transition between rolling contact (where the contact point has zero velocity relative to the terrain) and slipping is a crucial concept. Allowable traction force for soil is given as (Terzaghi, 1943):

$$\|F_{Tmax}\| = cA + \|F_N\| \tan \phi \quad (2)$$

where F_{Tmax} is the maximum force in the tangent direction, c is the soil cohesion, A is the contact area, F_N is the normal force and ϕ is the soil friction angle. The contact area A is estimated based on the sinkage due to normal force and the geometry of the wheel. Equation 2 shows that the normal force should be computed first since its magnitude limits the maximum traction at the contact. Figure 5 shows an overhead view of the tangent plane compliance system.

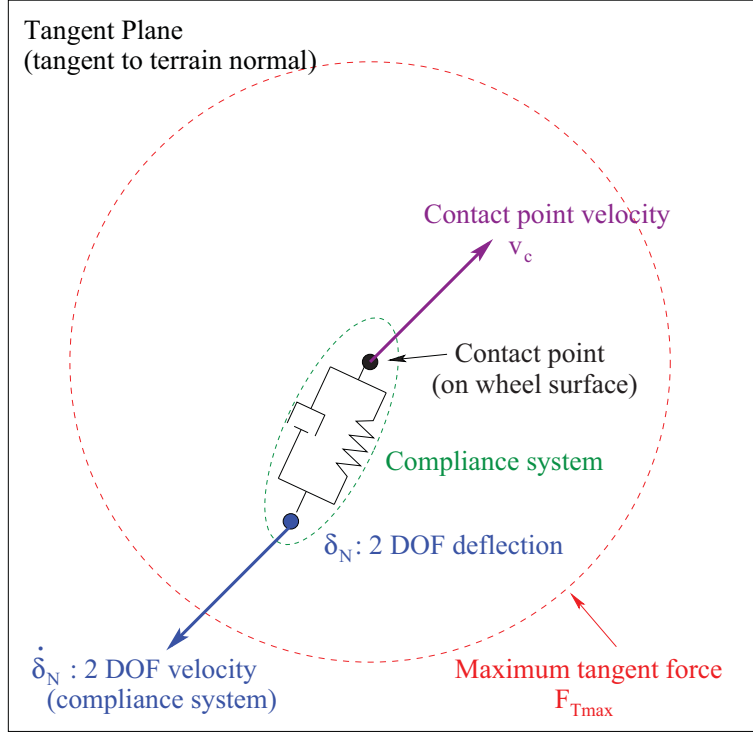


Figure 5: Tangent plane compliance system

In the **slipping regime**, the tangent plane force has reached its maximum allowable value (corresponding to the red, F_{Tmax} circle in Figure 5). The tangent plane force opposes the relative motion of the contact point. The tangent force is pointed in a direction opposite the relative velocity vector of the contact point. If the velocity of the contact point is not in the tangent plane, it will be projected into the tangent plane to determine the direction of the tangent force vector. The magnitude of the tangent force in the slipping regime is equal to the maximum allowable force:

$$\vec{F}_T = \|F_{Tmax}\| \frac{-\vec{v}_c}{\|\vec{v}_c\|} \quad (3)$$

where \vec{F}_T is the 2D tangent force and \vec{v}_c is the velocity of the contact point relative to the terrain projected onto the 2D tangent plane.

Computing the tangent plane force in the **rolling regime** is more problematic. In this regime, the velocity of the contact point with respect to the terrain is ideally zero. However, this does not mean that the tangent plane forces are zero. The only constraint is that the tangent plane force magnitude remain less than the maximum:

$$\|\vec{F}_T\| \leq \|F_{Tmax}\| \quad (4)$$

To solve the problem of computing tangent forces for the rolling regime, ROAMS uses a two degree of freedom compliance system as outlined in (Kraus et al., 1997). A linear spring-damper model is used:

$$\vec{F}_T = k_T \vec{\delta}_T + d_T \dot{\vec{\delta}}_T \quad (5)$$

where k_T is the spring coefficient, d_T is the damping coefficients and $\vec{\delta}_T$ is the 2D deflection in the tangent plane. The two degrees of freedom allow the compliance system to move in any tangent plane direction. Unlike the normal force compliant system, whose deflection is based on the position of the wheel and the local terrain height, the 2D deflection of the tangent plane system ($\vec{\delta}_T$) is not based on the rover state. Instead, the *derivative* of the deflection is based on the system state. Specifically, the derivative of the deflection is defined as the opposite of the velocity of the contact point relative to the surface:

$$\dot{\vec{\delta}}_T = -\vec{v}_c \quad (6)$$

Additional system states are added and used to track the actual deflection of the tangent plane compliance systems.

It is important to understand why the *derivative* of the deflection is defined in this manner. First, consider what it means for the velocity of the contact point relative to the surface to be zero. This velocity is zero when the wheel is in the rolling regime. If a torque is applied to a wheel resting on the surface, the contact point will begin to move relative to the surface in the absence of tangent plane forces. To oppose this motion, the tangent plane compliance system will deflect in the opposite direction - opposing the relative motion of the contact point. The resulting tangent plane forces propel the vehicle forward. The tangent plane compliance system acts as a controller that drives the tangent force to the correct rolling force. A small amount of slipping is accepted as this controller converges to the correct tangent force. The spring-damper coefficients used for the tangent plane compliance system serve as gains for this control system. Any torque applied to the wheel will change the equilibrium point for the control system and the tangent plane compliance system will attempt to deflect to a new point. The maximum traction force defines a limit on the deflection of the tangent plane compliance system. Once this deflection is reached, the wheel enters the slipping regime and no further deflection of the tangent spring is allowed.

Since the type of contact (slipping or rolling) is not known a-priori, ROAMS uses the following procedure to compute the tangent plane forces:

1. Assume contact is rolling and compute \vec{F}_T as

$$\vec{F}_T = k_T \vec{\delta}_T + d_T \dot{\vec{\delta}}_T \quad (7)$$

2. If $\|\vec{F}_T\| > \|F_{Tmax}\|$, contact is sliding:

$$\vec{F}_T = \|F_{Tmax}\| \frac{-\vec{v}_c}{\|\vec{v}_c\|} \quad (8)$$

The derivative of the tangent deflection, $\dot{\vec{\delta}}_T$ is computed by inverting Equation 7:

$$\dot{\vec{\delta}}_T = F_T^{-1}(\vec{\delta}_T, k_T, d_T) \quad (9)$$

3. If $\|\vec{F}_T\| \leq \|F_{Tmax}\|$, contact is rolling:

$$\dot{\vec{\delta}}_T = -\vec{v}_c \quad (10)$$

3.6 Soil randomness

The contact model described in 3.4 and 3.5 was originally formulated for rigid body contact. For two rigid bodies in contact, the transition between rolling and slipping regimes is very abrupt. Wheel-terrain contacts usually transition more gradually between regimes and can experience rapid fluctuations between rolling and slipping. Some slippage can occur even on perfectly flat terrain. These fluctuations are due to a wide range of effects that are not part of a simple rigid-body contact model, including variations in local soil parameters, soil deformation, wheel surface irregularities, vibratory effects, *etc.* While it is tempting to include more and more physical properties in the contact model, such modeling efforts are unrealistic for use in a general, real-time simulation tool as opposed to a detailed terra-mechanics analysis tool.

Equation 2 defines the tangent force where the contact transitions between rolling and slipping behavior. The contact model in ROAMS adds a Gaussian scaling factor to this equation:

$$||F_{Tmax}|| = s(cA + ||F_N|| \tan \phi) \quad (11)$$

where s is a Gaussian curve centered about a value of 1. This scale factor serves to “soften” the abrupt transition between rolling and sliding and more closely mimic the behavior of wheel-terrain contact. Some slipping may occur even on flat terrain (s is very small) and some rolling may occur on steep terrain (s is very large). The Gaussian scale factor s is defined by two parameters in ROAMS. The first parameter defines the maximum standard deviation of the Gaussian function and the second parameter defines how the standard deviation varies with local terrain angle. The standard deviation (s_{std}) for the scale factor s is:

$$s_{std} = p_1 - p_2 R \quad (12)$$

$$R = \begin{cases} \frac{\theta}{\phi} & , \quad \theta \leq \phi \\ \frac{\phi}{\theta} & , \quad \theta > \phi \end{cases} \quad (13)$$

where p_1 and p_2 are the two soil randomness parameters, ϕ is the soil friction angle and θ is the angle between the terrain normal and the vertical. The first parameter (p_1) defines the maximum standard deviation for the random Gaussian function. The second parameter (p_2) defines how the standard deviation varies with soil angle. The standard deviation is minimum when the terrain normal under a given wheel is equal to the soil friction angle ($R \rightarrow 1$). The standard deviation increases as the terrain normal becomes either much smaller or much larger than the friction angle ($R \rightarrow 0$). This increased standard deviation gives a chance of slipping for situations where rolling is the nominal behavior and vice versa. The choice of parameters p_1 and p_2 allows us to tune the behavior of the simulation and create a gradual transition between the rolling and slipping regimes.

The two following sections describe how the vehicle-terrain interaction model in ROAMS has been validated: (1) Validation based on Mars Exploration Rover experience (Section 4), and Validation based on Experiments with the FIDO rover (Section 5).

4 Validation based on Mars Exploration Rover (MER) Test Data

In order for a rover simulation to be useful in developing rover navigation and control software, its behavior must correspond well with the operation of the actual rover in a real

environment. Hence, in parallel with the ongoing development of ROAMS, we have undertaken a validation effort for ROAMS using experimental data from rover mobility runs. Rover motion is a product of many different components and levels of the system. At the lowest level, there are rover suspension components (rockers and bogeys), wheels and motors. While these low level rover components are usually well defined, the terrain used for experimental data collection is not. The terrain shape (height field) and physical properties (density, cohesion, friction angle) are important simulation parameters that are not typically well known for natural terrain used during rover mobility tests. The experimental mobility tests can also show wide variations between runs for the same rover over the same nominal terrain. The non-repeatable nature of these tests requires that the simulation use statistical techniques for comparison with empirical test data.

4.1 Parametric simulation

The wheel-soil contact model in ROAMS has a variety of parameters. While some of these parameters are physical (soil density, cohesion, friction angle), others are heuristic (soil randomness parameters p_1 and p_2). We use Monte Carlo testing to determine a set of model parameters that most closely matches experimental data from rover mobility testing. ROAMS has a framework to automate the collection of mobility data for a range of independent parameter sets. This data collection can be run in parallel on several computers to speed up the simulated data collection. A variety of motion primitives (straight, arc-turn, turn-in-place, *etc.*) on a variety of terrain types (flat, fixed slope, *etc.*) can be tested. The ROAMS mobility data is then compared to empirical data and the set of parameters that best matches the experimental data can be determined.

One example of tuning ROAMS parameters was the comparison of ROAMS against mobility testing on the Dynamic Test Model (DTM) of the Mars Exploration Rover (MER) rover. The DTM was designed to mimic the loads experienced by the MER rovers on the surface of Mars. The DTM rover is an earth-based testbed configured to have the same center of mass as the MER rover and approximately 120% of the MER rover’s Mars weight. Several mobility tests were conducted on Earth using the DTM rover to determine slippage when driving on slopes of up to 20 degrees. Figure 6 shows slip data collected during uphill, downhill and cross-hill mobility testing of the MER DTM rover on flat terrain at slope angles of 0, 2.5, 5, 10, 15 and 20 degrees (Lindemann, 2005). Figure 7 shows simulated results for straight uphill, straight downhill and straight cross-hill testing in ROAMS at 0, 5, 10, 15 and 20 degree slope angles. The ROAMS parameters were determined using Monte Carlo testing to find a parameter set that best matched the empirical data in Figure 6. The data used for comparison are shown as red circles in Figure 7. Since the empirical data does not have any statistical variance information, each data point was given equal weighting in determining the best set of ROAMS parameters. The resulting simulation shows a good match for cross-hill traversals, but the same parameter set did not provide as good a match for downhill motion (Figure 7). A technique for switching between parameters that best match empirical data in a given regime is being examined for use in ROAMS. This technique would allow ROAMS to modify parameters on-the-fly as the rover encounters different driving conditions.

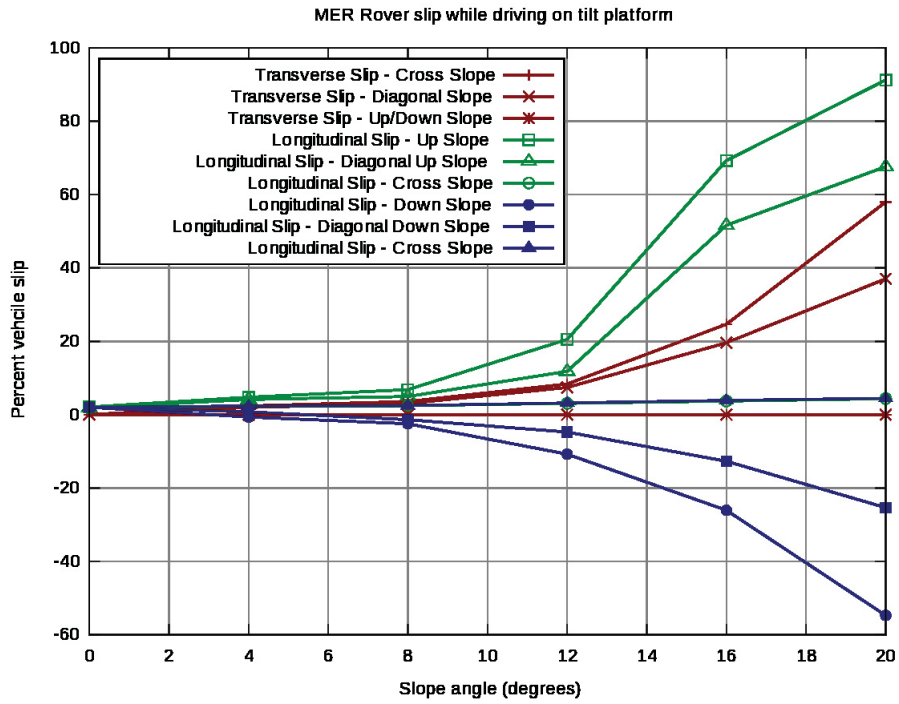


Figure 6: MER DTM rover test slip data

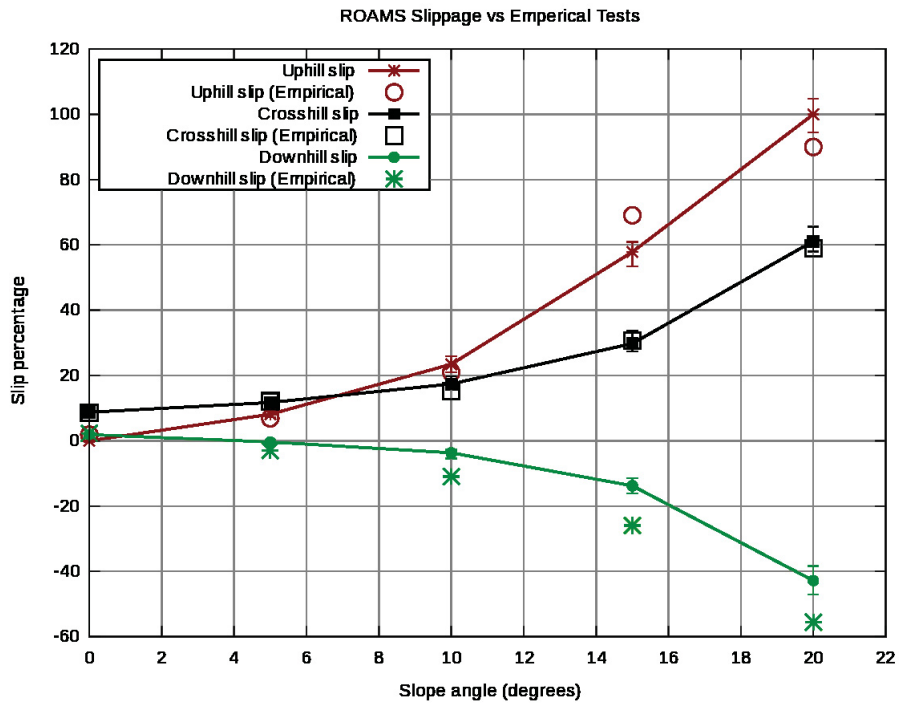


Figure 7: ROAMS straight uphill (with variance). Red circles are empirical data

5 Validation experiments with the FIDO rover

In order to tune and improve the behavior of ROAMS , a series of experimental tests was performed at JPL with the FIDO rover. The goals were to tune the vehicle-terrain interaction models to closely match those of the FIDO experimental rover. Some of the experimental tests were used to tune the ROAMS models and were then applied to the balance of the tests to analyze the performance of the ROAMS model compared to an experimental vehicle. This will be describe more detail now.

One of the key questions that must be addressed for acceptance of ROAMS for mission development and operational testing is how well does it represent reality. There are several aspects to this question. The first is whether the simulation captures the basic physics involved adequately. Assuming the first question is satisfied, the next question is whether the parameters involved in the simulation are tuned in such a way that the behavior of the simulation matches that of a specific real vehicle. Failure in the tuning is a clear indication that the basic models are inadequate. The goal in this activity was to tune the ROAMS simulation to the FIDO (Field Integrated Design and Operations) rover and show that both of these issues had been addressed. The FIDO rover is a six wheeled technology prototype with a rocker-bogie suspension system at JPL that was used for MER field trials throughout the period from 1999 to 2002 to develop operational strategies for mobility and science operations [4].

5.0.1 Identifying key ROAMS parameters

During development of the ROAMS simulation, significant work has gone into developing quality models of all aspects of rover kinematics and dynamics. In particular, special attention has been paid to modeling the interaction of the rover wheels with the terrain. Based on the resulting models and other consideration, there is a set of parameters that most significantly affect the ability of ROAMS simulation to match the behavior of the FIDO experimental runs. These parameters are:

- friction angle (radians),
- soil cohesion (kPa),
- rolling resistance (unitless),
- maximum standard deviation of slip transition parameter (for tire-soil slip model),
- slope standard deviation of slip transition parameter (for tire-soil slip model).

The first three parameters are typical ground trafficability parameters. The last two parameters are internal parameters in the wheel-soil interaction model relating to the stochastic nature of slip. Based on previous experience, nominal values of these parameters that worked well in simulations were known. One of the goals of this experiment was to also to identify the range of slopes that a single set of interaction parameters would be valid for. The results described here hence reflect the accuracy in matching experimental data with a single such parameter set.

5.0.2 Tuning to a subset of experimental runs

In order to tune these parameters, a simulation of the FIDO rover was developed with the ability to execute multiple simulation runs that sweep the set of parameters. These runs are referred to as “Monte-Carlo” runs, even though the selection of which parameters to use in each case was selected from a list of parameters instead of being chosen randomly, as in a true Monte-Carlo simulation. The fidelity of each simulation run was evaluated by constructing a score based on differences from the final position and heading in the simulation versus the experimental runs. A subset of the FIDO experimental runs was used and parametric sweeps runs for each test case were set up. Based on which simulations runs best matched the data from the actual runs, the simulation parameters were tuned to give the best overall match. The tuning was performed in two phases: a coarse phase and a refinement phase. In the first phase, a wider selection of values was used for each parameter about a nominal value guess. Based on comparing the parametric sweeps with the test cases, a new nominal value was selected that was approximately tuned to the FIDO performance. A second phase of tuning was then performed where a smaller range of variation of each parameter with respect to the new nominal values selected in the coarse tuning phase was considered. A wide range of parameters worked on runs for flat terrain, but a much smaller set produced good results on slopes. Fortunately, tuning for slopes produced good results on the flat. In general, it was more difficult to match the behavior of the real FIDO rover on steeper slopes (greater than approximately 15 degrees).

5.1 Experimental studies

The characterization of the ROAMS simulation technology with the FIDO rover was conducted from August 31, 2004 to September 23, 2004 using the MER tilt table adjacent to the JPL MarsYard. The tilt table has a mosaic of paving stones glued to a plywood surface and covered with a layer of fine sand (shown in Figures 8 and 9). The surface setup of the tilt table was used by the MER engineering flight team to determine the operating characteristics of the Opportunity rover as it descended towards the crater floor in Endurance Crater at Meridiani Planum. Driving strategies were derived from those engineering tests and then used on Mars.

The ROAMS parameter settings were obtained by using a subset of the test data to tune the ROAMS model for FIDO that then served as the baseline for simulation performance evaluation. Tests were run on the tilt table set at 0, 5, 10, and 20 degrees. The actual tilt table angles used for the tests measured with a Leica TotalStation were 0.085, 5.743, 10.758, and 20.764 degrees, but these will be referred to as 0, 5, 10, and 20 degrees throughout this section. A total of 400 tests were run to determine ROAMS simulation performance for accuracy during straight-line and arc drives, and turns-in-place at the 4 tilt table angles.

All results were compared to ground truth data that was collected by a Leica TCRA1103+ TotalStation (rated accuracy of 2mm at 10m range). All of the test runs were also taped on Mini-DV video for possible future analysis of the full behavior of the rover. The test setup is shown in Figures 10 and 11, where the FIDO rover has four prism reflectors (one on each corner of the rover) that are used for range and rover attitude determination each time



Figure 8: Underlying surface of the MER tilt table covered with concrete paving stones glued to a plywood surface.



Figure 9: The MER tilt table covered with sand for experiments.



Figure 10: Leica TotalStation used in the FIDO experimental setup



Figure 11: FIDO Rover on experimental slope table

before and after the tests. This setup gives both position and heading information. All tests were run using the current version of the FIDO control code. Logfiles were produced by the FIDO control software at a rate of 50Hz and were used for ROAMS parameter setting and characterization runs.

5.1.1 Drive runs

There were a total of 88 different types of drive tests run with the tilt table set at 0, 5, 10, and 20 degrees. Each one of the tests was run five times for the level table and three times for the tilted table settings (5, 10, and 20 degrees). The drive tests included straight line, shallow arcs, and tight arcs of length 2 meters. These drives were done with up-slope, down-slope, and cross-slope directions and the rover traveling forward and backward. The combination of rover driving direction, slope direction, and type of drive (straight vs arc) fully covers the types of drives and environmental conditions experienced by planetary rovers such as MER.

The average heading error (in degrees) defined as the difference between the ground truth and the ROAMS outputs for each tilt table setting (0, 5, 10, and 20 degrees) for all of the drive runs is shown in Figure 15. The largest errors are seen in the tight arcs due to interaction of the wheels with the underlying mixed sand and concrete paving stones (see Figure 8 and 9). This behavior also leads to the larger standard deviations for these runs (see Tests 19 and 20 in Figure 15) due to the extreme sensitivity of small variations in heading during a tight arc. Also of particular importance for evaluation of the fidelity of ROAMS is the slip estimate while driving on slopes. A scatter plot of the measured slip versus ROAMS slip (in percent) based on the commanded final position for each drive run is shown in Figure 12.

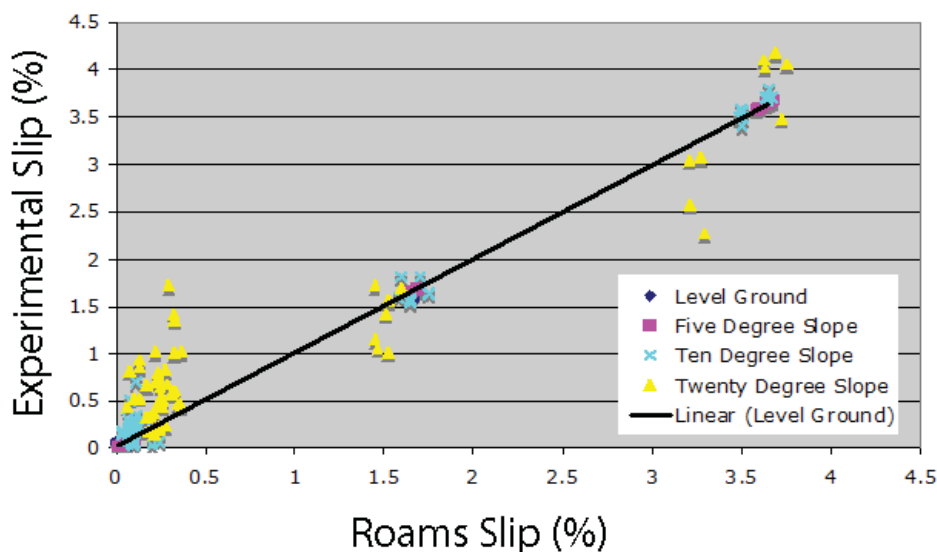


Figure 12: Plot of measured slip versus ROAMS simulation slip for all MER tilt table angles (0, 5, 10, and 20 degrees). Note the wider variance in the slip for the 20 degree slope.

The variance in slip estimates grows with MER tilt table angle, especially in the lower slip range for the 20-degree setting, where ROAMS consistently underestimates the slip. Examples of ROAMS output plots for one of the some typical runs is shown in Figures 13 and 14.

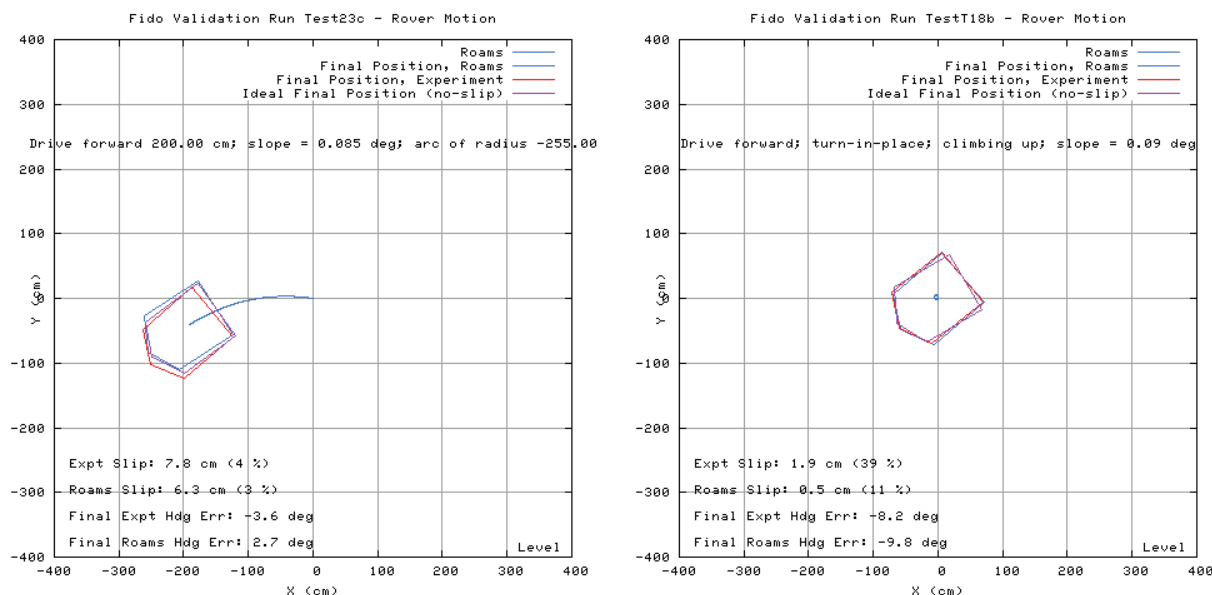


Figure 13: Typical ROAMS and experimental results for FIDO tests on a level surface. Blue represents ROAMS results, experimental results are shown in red, and the ideal final results are shown in purple. The figure on the left is for driving an arc cross-slope and the figure on the right shows a turn in place run. Most tests on level or 5-degree slopes show similar excellent agreement.

The following table summarizes the percent slip (as compared to the commanded motion) for all of the FIDO experiments compared to ROAMS. (A few outliers were omitted for FIDO experimental runs outside nominal rover behavior.)

Slope	FIDO Experiment % slip		FIDO ROAMS % slip	
	Min	Max	Min	Max
Level	2	5	2	4
5 deg	2	8	1	6
10 deg	2	20	2	10
20 deg	10	60	5	20

Table 1: Summary of slip differences between the FIDO Experimental data and ROAMS for all runs

The summary graphs indicate that the average error in heading for the turn-in-place runs were usually less than ten percent for slopes less than twenty degrees. In seven of the eighteen

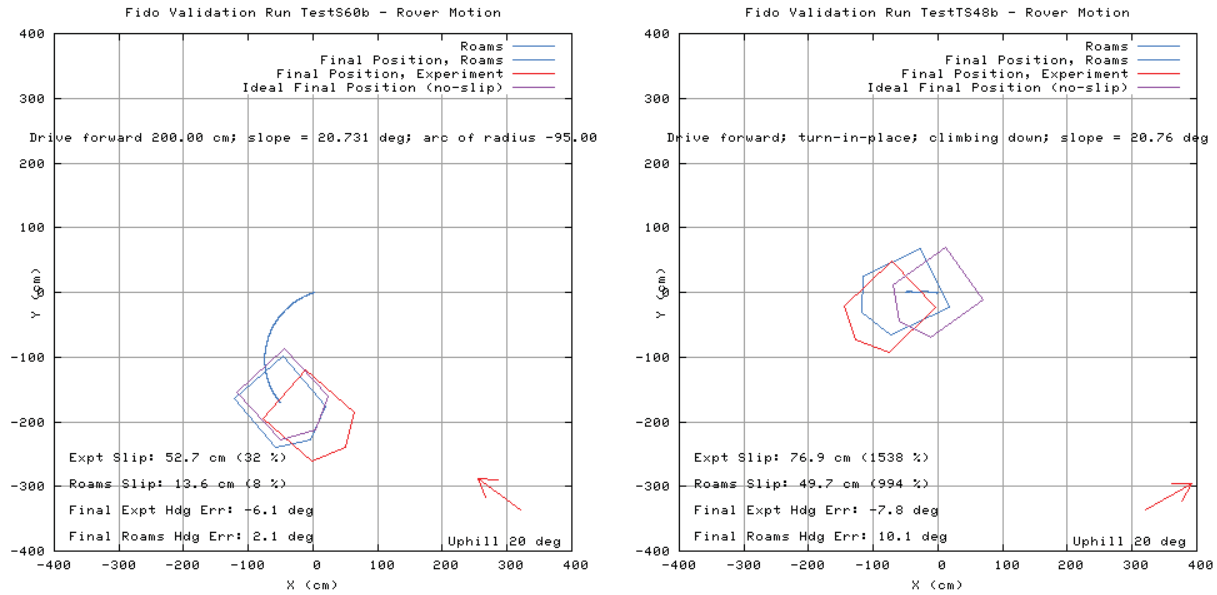


Figure 14: Typical ROAMS and experimental results for FIDO tests on a 20-degree slope. Blue represents ROAMS results, experimental results are shown in red, and the ideal final results are shown in purple. The figure on the left is for driving an arc cross-slope and the figure on the right shows a turn in place run. The uphill direction is shown with the arrow in the bottom right corner of the figure. There was extensive down-slope slip on 20 degree runs.

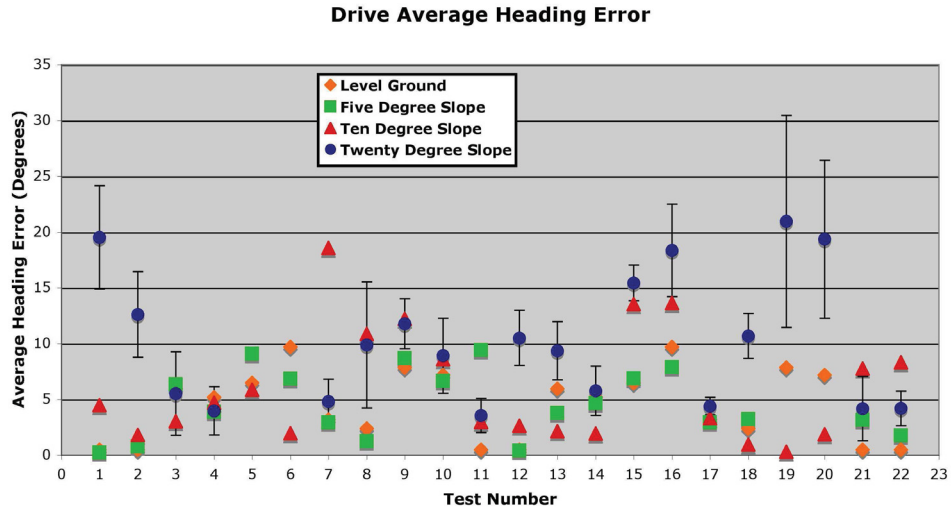


Figure 15: Summary plot of heading errors for the drive runs on the MER tilt table set at 0, 5, 10, and 20 degrees. The heading error here is dened as absolute difference between the ground truth and the ROAMS output. Standard deviation error bars are shown only on the 20 degree slope data values for clarity of presentation.

Tests run at a twenty degree slope setting, the average errors were greater than ten percent and as high as thirty-four percent for one of the Tests. The summary graphs indicate that

the average error in heading for the drive runs were usually less than six degrees for slopes less than twenty degrees. In twelve of the twenty runs at a twenty degree slope setting, the average errors were greater than six degrees and as high as twenty-one degrees for one of the runs.

In terms of localization in the X and Y coordinate frame, the majority of the error was in the X direction (aligned with the forward direction of the rover). The majority of the errors were less than ten centimeters in the X and Y directions for slopes less than ten degrees, growing to twenty centimeters in the X and Y directions for slopes of ten degrees, and increasing to forty centimeters in the X and Y directions for slopes of twenty degrees. These errors are all measured for two meter traverses.

5.1.2 Turn-in-place runs

There were a total of 18 types of runs for the turn-in-place tests with tilt table set at 0, 5, 10, and 20 degrees. Each one of the tests was run five times for the level table and three times for the tilted table settings (5, 10, and 20 degrees). The turn-in-place runs were done with right and left turns of 45, 90, and 360 degrees in upslope, down-slope, and cross-slope directions and the rover traveling forward and backward. The average heading error (in percent) defined as the difference between the ground truth commanded angle and the ROAMS outputs for each tilt table setting (0, 5, 10, and 20 degrees) for all of the turn-in-place runs is shown in Figure 16. Of particular note is the increase in the average heading errors when the MER tilt table was set to 20 degrees and the rover was facing downhill (Tests 7 and 8) where the standard deviation is larger due to the rover having to fight the gravity vector combined with snagging on the underlying paving stones..

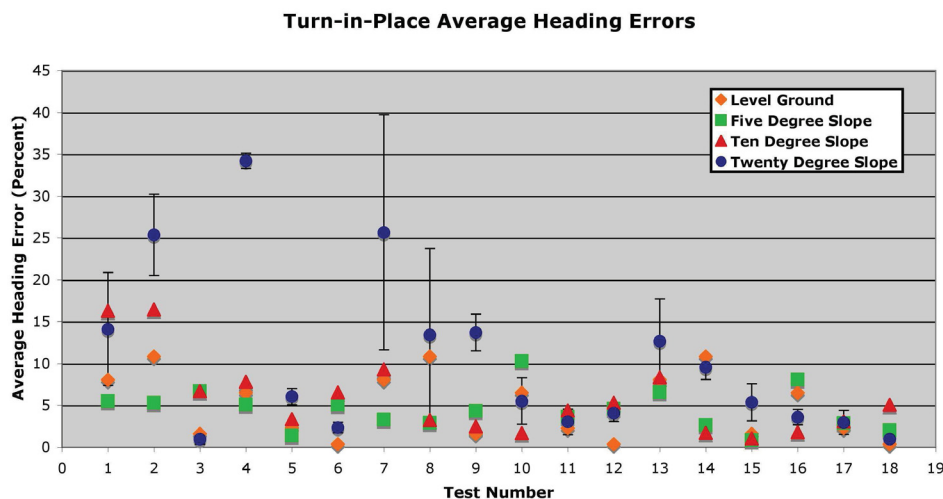


Figure 16: Summary plot for heading error for the turn-in-place runs on the MER tilt table set at 0, 5, 10, and 20 degrees. The heading error here is dened as the percentage difference between the ground truth commanded angle to turn and the ROAMS output. Standard deviation error bars are shown only on the 20 degree slope data values for clarity of presentation.

6 Summary and future directions

The overall results indicate that in terms of prediction error, ROAMS generally simulated the field collected experimental data well for MER tilt table for slopes up to and including ten degrees. There is a large amount of variability in the FIDO experimental data for the twenty degree slope setting on the MER tilt table. Generally, one out of three test runs would be outside the nominal behavior of the rover. The factors that contribute to this behavior are the excessive slip at this slope angle, interactions with the underlying paving stones, and bogie interactions with the terrain. The large amount, and variability, in slippage (10% to 60%) observed at twenty degrees slope was sometimes underestimated by ROAMS .

In conclusion, excellent validation of ROAMS vehicle-terrain interaction has been demonstrated for level to medium slopes (matching the measured slip to within 1 to 2 percent). For steeper slopes, the ROAMS simulation did not match the experience of the test vehicles as well (sometimes off by a factor of 2 in the percentage slip). But it is also clear that the motion of the experimental vehicles is somewhat chaotic on steep slopes and therefore more difficult to match in simulations. While ROAMS captured the chaotic behavior on steep slopes, the deviations between the slippage predictions and actual data were large. This work did not focus on trying to match the motion distributions on large slopes. The majority of the slopes experienced by the MER Spirit and Opportunity have been below 20 degrees, so the expectation is that the bounds on the accuracy of the ROAMS characterization described in this paper will cover the usual driving circumstances experienced on Mars. In the event that steeper terrain is encountered, the hardware testbed sandbox such as that used for MER would be employed. There is room for improvement in ROAMS vehicle-terrain modeling, especially for steeper slopes, but ROAMS is already a useful tool for performing realistic vehicle simulations of driving on Mars as was done in the SOOPS project (Haldeman et al., 2007).

An obvious next step to improve slippage modeling for large slopes is to vary the wheel/soil interaction model parameters across different slope regimes instead of limiting the model to a single parameter set for all slopes. This would allow improved matching of motion distributions on large slopes with the large distributions of the experimental vehicles. Future improvements to the validation experiments include the measurement and use of variable terramechanics properties instead of uniformly constant values. Extending the experiment to different terrain types and a larger terrain set is also highly desirable. A probabilistic component will be added to the ROAMS simulation model in order to handle the off-nominal behavior of rovers on slopes steeper than ten degrees (the MER rover Spirit has driven on slopes greater than twenty-five degrees and the Opportunity rover recently performed a carefully planned drive on a twenty-eight degree slope below the lip of Victoria Crater). Further research will investigate the use of more sophisticated soil contact models using finite element algorithms that relax the current single point contact assumption.

The validation results described in this paper are specifically FIDO and MER rovers, representing a particular class and size of rovers. Another area for future development is to apply similar validation approaches to larger vehicles. For instance, the rover planned for MSL , the next Mars mission, is the same general type of rover as FIDO or MER but it is several

times larger and heavier. Further validation work is needed to determine how well current models will scale up from FIDO to MSL. Validation work would also be beneficial to other types of robotic vehicles, such as ATHLETE (a 6-legged vehicle with wheels at the end of each leg proposed for use on the moon) and CHARIOT (a large flat-bed rover planned for use by astronaut transportation on the moon).

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

- ADAMS (2000). *ADAMS User Manual*. Dynamic Computation Software, Inc.
- Ahlvin, R. B. and Haley, P. W. (1992). Nato reference mobility model, edition ii, nrmm ii user’s guide. In *Technical report GL-92-19, (AD 170301, US Army Engineer Waterways Experiment Station, Vicksburg, MS*.
- Alexander, R. K. (1994). Stability of Runge-Kutta methods for stiff ordinary differential equations. *SIAM Journal on Numerical Analysis*, 31(4):1147–1168. (cite-seer.ist.psu.edu/alexander94stability.html).
- Andrade, G., Amar, F. B., Bidaud, P., Chatila, R., and Velizy, L. (1998). Modeling robot-soil interaction for planetary rover motion control. In *Proceedings 1998 International Conference on Intelligent Robots and Systems*, volume 1, pages 576–581.
- Bauer, R., Leung, W., and Barfoot, T. (2005). Development of a dynamic simulation tool for the exomars rover. In *Proceedings 2005 International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Munich, Germany.
- Bekker, M. (1969). *Introduction to Terrain-Vehicle Systems*. University of Michigan Press.
- Bekker, M. G. (1956). *Theory of Land Locomotion: The Mechanics of Vehicle Mobility*. University of Michigan Press.
- Bekker, M. G. (1960). *Off-the-Road Locomotion: Research and Development in Terramechanics*. University of Michigan Press.
- Biesiadecki, J., Henriquez, D., and Jain, A. (1999). A reusable, real-time spacecraft dynamics simulator. In *Proceedings of 6th Digital Avionics System Conference*, Irvine, CA.
- Chan, B., Thompson, N., Sandu, C., and Dankowicz, H. (2005). A novel tire-soil interaction model for simulation of off-road vehicles. In *Proceedings of the 15th Conference of the ISTVS*, Hayama, Japan.
- D. J. van Wyk, J. Spoelstra, J. H. d. K. (1996). Mathematical modeling of the interaction

- between a tracked vehicle and the terrain. *Applied Mathematical Modeling*, 20(11):838–846.
- Ellery, A. (2005). Environmentrobot interactionthe basis for mobility in planetary micro-rovers. *Robotics and Autonomous Systems*, 51(1):29–39.
- Farritor, S., Hacot, H., and Dubowsky, S. (1998). Physics-based planning for planetary exploration. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, volume 1, pages 278–283, Leuven, Belgium. (citeseer.ist.psu.edu/396131.html).
- Hacot, H. (1998). *Analysis and Traction Control of a Rocker-Bogie Planetary Rover*. Massachusetts Institute of Technology. M.S. Thesis.
- Hacot, H., Dubowsky, S., and Bidaud, P. (1998). Analysis and simulation of a rocker-bogie exploration rover. In *Proceedings of the Twelfth CISM-IFTMM Symposium on Theory and Practice of Robots and Manipulators*, Paris, France. (citeseer.ist.psu.edu/398566.html).
- Hait, A. and Siméon, T. (1996). Motion planning on rough terrain for an articulated vehicle in presence of uncertainties. In *Proceedings of 1996 IEEE International Conference on Robots and Systems*, Osaka, Japan. (citeseer.ist.psu.edu/hait96motion.html).
- Haldeman, A., McHenry, M., Petras, R., Bornstein, B., Castrano, R., Cameron, J., Estelin, T., Farr, T., Gaines, D., Jain, A., Leff, C., Lim, C., Nesnas, I., Pomerantz, M., Powell, M., Shu, I., and Volpe, R. (2007). Simulation to evaluate autonomous behaviors for mobile planetary surface science missions. In *Proceedings 2007 IEEE Aerospace Conference*, Big Sky, Montana.
- Haley, P. W., Jurkat, M. P., and P. M. Brady, J. (1979a). Nato reference mobility model, edition i, users guide, volume i, operational modules. In *Technical Report 12503, (AD B047979L)*, US Army Tank-Automotive Research and Development Command, Warren, MI.
- Haley, P. W., Jurkat, M. P., and P. M. Brady, J. (1979b). Nato reference mobility model, edition i, users guide, volume ii, obstacle modules. In *Technical Report 12503, (AD B047980L)*, US Army Tank-Automotive Research and Development Command, Warren, MI.
- Hunt, K. and Crossley, F. (1975). Coefficient of restitution interpreted as damping in vibroimpact. *Journal of Applied Mechanics*, pages 440–445.
- Huntsberger, T., Jain, A., Cameron, J., Woodward, G., Meyers, D., and Sohl, G. (2008). Characterization of the roams simulation environment for testing rover mobility on sloped terrain. In *Proceedings of 9th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*, Los Angeles, CA.
- Hutangkabodee, S., Zweiri, Y., Seneviratne, L., and Althoefer, K. (2008). Soil parameter identification and driving force prediction for wheel-terrain interaction. *International Journal of Advanced Robotic Systems*, 5(4):425–432.
- Iagnemma, K., Genot, F., and Dubowsky, S. (1999). Rapid physics-based rough-terrain rover planning with sensor and control uncertainty. In *Proceedings of 1999 IEEE International Conference on Robotics and Automation*, pages 2286–2291, Detroit, MI. (citeseer.ist.psu.edu/iagnemma99rapid.html).

- Jain, A., Balaram, J., Cameron, J., Guineau, J., Lim, C., Pomerantz, M., and Sohl, G. (2004). Recent developments in the roams planetary rover simulation environment. In *Proceedings of 2004 IEEE Aerospace Conference*, Big Sky, Montana.
- Jain, A., Cameron, J., Lim, C., and Guineau, J. (2006). Simscape terrain modeling toolkit. In *Proceedings of Second International Conference on Space Mission Challenges for Information Technology (SMC-IT 2006)*, Pasadena, USA.
- Jain, A., Guineau, J., Lim, C., Lincoln, W., Pomerantz, M., and Sohl, G. (2003). Roams: Planetary surface rover simulation environment. In *Proceedings of 2003 International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*, Nara, Japan.
- Jain, A. and Yen, J. (1999). Roams: Rover analysis modeling and simulation software. In *Proceedings of 1999 International Symposium on Artificial Intelligence, Robotics, and Automation in Space (iSAIRAS)*, Noordwijk, Netherlands.
- Kewlani, G. and Iagnemma, K. (2008). A stochastic response surface approach to statistical prediction of mobile robot mobility. In *Proceedings 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nice, France.
- Kraus, P., Fredriksson, A., and Kumar, V. (1997). Modeling of frictional contacts for dynamic simulation. In *Proceedings of IROS 1997 Workshop on Dynamic Simulation: Methods and Applications*. (citeseer.ist.psu.edu/kraus97modeling.html).
- Lindemann, R. (2005). Dynamic testing and simulation of the mars exploration rover. In *Proceedings of ASME International Design Engineering Technical Conference*, Long Beach, CA.
- Pacejka, H. B. and Bakker, E. (1992). The magic formula tyre model. *Vehicle Systems Dynamics*, 21(1):1–18.
- Poulakis, P., Joudrier, L., Wailliez, S., and Kapellos, K. (2008). 3drov: A planetary rover system design, simulation and verification tool. In *Proceedings 9th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Los Angeles, CA.
- Q. Li, P. D. Ayers, A. B. A. (2007). Modeling of terrain impact caused by tracked vehicles. *Journal of Terramechanics*, 44(6):395–410.
- Rodriguez, G., Jain, A., and Kreutz-Delgado, K. (1992). Spatial operator algebra for multi-body system dynamics. *Journal of the Astronautical Sciences*, 40:27–50.
- Sandu, C., Sandu, A., and Li, L. (2006). Modeling of stochastic terrain profiles and soil parameters. *SAE 2005 Transactions Journal of Commercial Vehicles*, 114(2):221–230.
- Shibly, H., Iagnemma, K., and Dubowsky, S. (2005). An equivalent soil mechanics formulation for rigid wheels in deformable terrain, with application to planetary exploration rovers. *Journal of Terramechanics*, 42(1):1–13.
- Singh, S. (1997). The state of the art in automation of earthmoving. *ASCE Journal of Aerospace Engineering*, 10(4).
- Sohl, G. and Jain, A. (2005). Wheel-terrain contact modeling in the roams planetary rover simulation. In *Proceedings of ASME International Design Engineering Technical Conference*, Long Beach, USA.

- Solisa, J. M. and Longoria, R. G. (2008). Modeling trackterrain interaction for transient robotic vehicle maneuvers. *Journal of Terramechanics*, 45(3):65–78.
- Terzaghi, K. (1943). *Theoretical Soil Mechanics*. John Wiley.
- Thueer, T., Krebs, A., Siegwart, R., and Lamon, P. (2007). Performance comparison of rough-terrain robots - simulation and hardware. *Journal of Field Robotics*, 24(3):251–271.
- Ward, C. C. and Iagnemma, K. (2007). Model-based wheel slip detection for outdoor mobile robots. In *2007 IEEE International Conference on Robotics and Automation*, Roma, Italy.
- Wong, J. Y. (2001). *Theory of Ground Vehicles, Third Edition*. John Wiley and Sons.
- Wong, J. Y. and Huang, W. (2006). wheels vs. tracks a fundamental evaluation from the traction perspective. *Journal of Terramechanics*, 43(1):27–42.
- Yen, J. (2008). Rover slip validation and prediction for mars exploration rovers. *Sensors and Transducers Journal (special edition)*, 90:233–242.